arXiv:1702.07319v1 [q-bio.NC] 23 Feb 2017

Learning to Draw Dynamic Agent Goals with Generative Adversarial Networks

Shariq Iqbal SHARIQ.IQBAL@DUKE.EDU
Duke Institute for Brain Sciences,
308 Research Dr.,
Durham, NC 27708 USA

John Pearson JOHN.PEARSON@DUKE.EDU
Duke Institute for Brain Sciences,
308 Research Dr.,
Durham, NC 27708 USA

Abstract

We address the problem of designing artificial agents capable of reproducing human behavior in a competitive game involving dynamic control. Given data consisting of multiple realizations of inputs generated by pairs of interacting players, we model each agent's actions as governed by a time-varying latent goal state coupled to a control model. These goals, in turn, are described as stochastic processes evolving according to player-specific value functions depending on the current state of the game. We model these value functions using generative adversarial networks (GANs) and show that our GAN-based approach succeeds in producing sample gameplay that captures the rich dynamics of human agents. The latent goal dynamics inferred and generated by our model has applications to fields like neuroscience and animal behavior, where the underlying value functions themselves are of theoretical interest.

1. Introduction

While recent years have seen tremendous advances in the training of sophisticated game-playing agents using reinforcement learning methods Mnih et al. (2015); Silver et al. (2016), comparatively little effort has been directed at a related goal, that of designing artificial agents that *play like humans*. This second goal is itself interesting for a number of reasons, including the possibility that human players may find it more rewarding to play against artificial opponents that share their biases, and that much of the research on the psychology of decision-making is more concerned with understanding human choices and their motivation *per se* rather than whether such play is optimal. In these cases, we are therefore more interested in modeling (and usually generating) human-like behavior than in the classical reinforcement learning problem; instead of teaching an agent to solve the task, we are concerned with reverse-engineering the strategies an agent has learned to use.

In this work, we are particularly interested in the neuroscientific applications of such an algorithm. Nearly all neuroscience experiments in humans and non-human animals feature a task in which the subjects perform some behavior repeatedly, often while brain activity is recorded. These brain signals are then correlated with variables of interest as a means of testing hypotheses about the computations underlying the observed behavior. Unfortunately, for many decision processes, the variables of interest — value functions, control signals, goals — are latent, and must be inferred. When experimental designs are simple, simple models of behavior suffice, but when the range of allowed behaviors in the task is complex (e.g., multiplayer interaction with continuous inputs) models must be flexible enough to produce a credible account of behavior, or the inferred variables will be useless as signals of interest.

Within the decision science community, the literature on inferring beliefs and value functions from observed behavior is vast Camerer (2003). However, the games to which these techniques have been applied are relatively constrained, typically consisting of a discrete state space; sequential, not simultaneous movement;

and a low-dimensional set of allowed actions. Here, we focus on a somewhat different problem, of the sort more typical in machine learning challenges: we assume a continuous state space, controlled by continuous (e.g., joystick) inputs, and simultaneous movement by both players. This is much closer in spirit to the inverse reinforcement learning problem Ng et al. (2000); Abbeel and Ng (2004), which attempts to replicate behaviors that match exemplars generated by an expert demonstrator.

In this work, we propose a method that learns to produce such time series by modeling the dynamics of each player's goal state as a function of time. These goals are chosen at each instant based on both the current goal and state of the game and subsequently drive game input via a simple control model. The resulting value functions are both parsimonious as a description of the game's dynamics and useful as a means of understanding players' choices, since in cases where goals are onscreen locations, they can be directly visualized. Key to the model's ability to generate realistic trajectories is the use of generative adversarial networks or GANs Goodfellow et al. (2014) to approximate successive draws of the goal state. GANs have recently excelled at generating novel faces and scenes after training on image databases Goodfellow et al. (2014); Mirza and Osindero (2014); Radford et al. (2015), but applications to time series modeling have been rarer. Our results show that, just as with images, GAN methods succeed in capturing the multimodality apparent in real data.

The outline of the paper is as follows: In Section 2, we outline our assumptions about the data and control model. In Section 3, we describe our GAN-based model for the latent goal states. In Section 4 we describe our full generative model for the data along with an algorithm for performing approximate Bayesian inference on the unknown model parameters. In Section 5 we describe our training procedure and compare the behavior of our model with the dynamics exhibited by real players in the same task. We show that not only are we able to replicate the wide range of natural behavior, but that the inferred value functions provide an intuitive means of accounting for this variability. Section 6 references related work and details novel contributions of this work. Section 7 concludes and discusses potential applications of our model.

2. Data and control model

The aim of our model is to use a set of observed movement trajectories from each player in a two-player game to infer an underlying goal dynamics capable both of explaining behavior and simulating actions from a realistic artificial opponent. To do so, we combine a model of latent goal states (with dynamics generated by a value function) with a simple control model that translates these goals into actions. Since our model is generative, the final result is able to produce entirely new behavior that captures the variability present in real opponents.

2.1 Data

The inputs to our model are the complete time series for all relevant onscreen variables, y_t, and all control inputs by both players, u_t. We also consider a state variable for the system, s_t, consisting of all variables on which agents' behavior might be conditioned (y_t, velocities, score, etc.). We assume that the training set consists of a large number of bouts of play, with each bout constituting a single realization of each time series.[1] We further assume that control inputs are translated into onscreen variables via

$$y_t = y_{t-1} + v_{\max} \odot \sigma(u_t) \tag{1}$$

with $v_{\max}$ a scaling between control and velocity, $\sigma(x)$ a sigmoid function reflecting saturation of control signals, and $\odot$ elementwise multiplication.

1. In this work, we assume these realizations are iid, though future work might consider modeling changes in the value function of each agent during the course of play.

2.2 Control model

We assume that at each moment in time, each agent has a desired goal state g_t (often an onscreen position) and a controller capable of achieving this state by minimizing the error $e_t \equiv g_t - y_t$. For example, a typical choice is the proportional integral derivative (PID) controller, which takes the discrete form (for a single control variable)

$$\Delta u_t = \kappa \left[\alpha e_t - \beta e_{t-1} + \gamma e_{t-2}\right] \tag{2}$$

where κ is the proportional control constant and α, β, and γ are functions of T_i and T_d, the integration and differentiation time constants. More generally, we can write the change in control signal as a convolution: $\Delta u = L * (g - y)$, with L an *a priori* unknown linear control operation. In what follows, we will also assume some Gaussian uncertainty in the control signal

$$u_t \sim \mathcal{N}(u_{t-1} + L * (g_t - y_{t-1}), \epsilon^2) \tag{3}$$

while keeping the relationship (1) between u_t and y_t deterministic.

3. Goal model

For the goal time series, we will assume a Markov process in which new goals are probabilistically selected at each time based on both the current goal and the current state of the system. That is,

$$p(u, g) \propto \prod_t p(u_t | u_{t-1}, g_t, y_{t-1}) p(g_t | g_{t-1}, s_{t-1}) \tag{4}$$

More specifically, we will assume that at each time point, there exists a value function $V(g_t, s_{t-1})$ that captures the benefit in setting a particular goal at the next time step based on the current state of the system. That is, we want to increase V as often as possible[2]. However, we add as a regularization constraint the idea that there should be some cost to large changes in goals, which we take to be quadratic in the distance between successive points. Explicitly, let

$$\log p(u, g) = \sum_t \left[-\frac{1}{2\epsilon^2} \|u_t - u_{*t}(g_t, y_{t-1})\|^2 \; -\frac{1}{2\sigma^2} \|g_t - g_{t-1}\|^2 + V(g_t, s_{t-1}) \right] - \log Z \tag{5}$$

Here, $u_{*t} \equiv u_{t-1} + L*(g_t - y_{t-1})$ is the the predicted control and ϵ and σ govern the noise in the observations and goal diffusion, respectively. In what follows, we will also find it useful to define $\beta \equiv \sigma^{-2}$ in order to write $\log p(g) = -\beta E(g|s) - \log Z$ with

$$E(g|s) = \sum_t \left[\frac{1}{2} \|g_t - g_{t-1}\|^2 - \sigma^2 V(g_t, s_{t-1}) \right] \tag{6}$$

This formulation admits multiple interpretations: In analogy with the path integral formulation of stochastic processes, it can be viewed as a model in which the "potential energy" $U(g) \equiv -\sigma^2 V$ trades off at each time step with the "kinetic energy" $K \equiv \dot{g}^2/2$. In the limit of small V/small σ/large β, one is in either the low-temperature thermodynamic limit or the high-mass classical limit, and g is a spatially-varying perturbation of a Gaussian process. Alternately, in the limit of large σ, goals are simply chosen independently at each time point. In any case, we have made the strong assumption that dependence of g_t on g_{t-1} occurs only through a momentum term, requiring that the "static" V term carries most of the weight of explanation.

Unfortunately, for general $V(g|s)$, the distribution implied by (6) is of the Boltzmann-Gibbs form and is impossible to sample efficiently. And while it is in principle possible to train a generative network such as a GAN to sample innovations in the g time series directly, conditioning on previous goal, this intermixes the dynamics of the system with the value function. If the goal of our inference is to model V itself, we will need a method for sampling from $p(g)$ that still allows this partitioning.

2. Alternately, we could consider $-V$ as an energy function that we wish to minimize.

3.1 $V(g)$ as a mixture of Gaussians

However, for any V defining an absolutely continuous Boltzmann distribution it is possible to rewrite the potential energy piece of (6) as:

$$e^{V(g)} \propto \mathbb{E}_{\mu,\lambda}\left[e^{-\frac{\lambda}{2\sigma^2}\cdot(g-\mu)^2}\sqrt{\frac{\lambda}{2\pi\sigma^2}}\right] \tag{7}$$

That is, e^V is a mixture of Gaussians, with mixture proportions determined by some unknown $p(\mu, \lambda)$. Clearly, if we can efficiently sample (μ, λ), then sampling from e^V is simply a matter of drawing one such value and subsequently sampling g_t from

$$p(g_t|g_{t-1}, \mu_t, \lambda_t, \sigma^2) \propto \exp\left(-\beta\left[\frac{1}{2}\|g_t - g_{t-1}\|^2 + \frac{\lambda_t}{2}\|g_t - \mu_t\|^2\right]\right) \tag{8}$$

which is equivalent to

$$g_t|g_{t-1}, \lambda_t, \mu_t, \sigma^2 \sim \mathcal{N}\left(\frac{g_{t-1} + \lambda_t\mu_t}{1 + \lambda_t}, \frac{\sigma^2}{1 + \lambda_t}\right) \tag{9}$$

That is, conditioned on previous goal position, μ, and λ, the new goal is simply governed by a normal distribution. Thus, assuming we can sample (μ_t, λ_t), we have a means of sampling entire goal trajectories. In addition, in the limit of large β, V can be easily approximated by an importance sampling procedure (see Supplement).

Our approach to modeling the mixture probabilities $p(\mu, \lambda)$ is based on a conditional version of a generative adversarial network (GAN) Goodfellow et al. (2014); Mirza and Osindero (2014). Assuming a set of existing samples $\{(\mu, \lambda, s)\} \sim p(\mu, \lambda|s)$[3], this generator, typically a neural network, learns to draw approximate samples from p by competing in a minimax game against a discriminator network, the goal of which is to discern between real samples and those produced by the generator. More formally, given samples ξ from a known distribution (e.g., multivariate normal or uniform), the generator network transforms these to $(\mu, \lambda) = G(\xi|s)$, while the discriminator function D classifies them as either real or synthetic: $D(\mu, \lambda|s) \to \{0, 1\}$. Training proceeds by alternately changing D to maximize its performance on a fixed set of real and synthetic samples and altering G so that its new samples minimize the performance of D. Model training is complete when the optimized discriminator performs no better than chance in differentiating between real data and generated samples. In a range of challenging applications, GANs have produced state-of-the-art performance, allowing for efficient sampling of complex probability distributions without explicit modeling of their form.

In general, training GANs requires some care. We use the recently proposed Wasserstein GAN (WGAN) training objective Arjovsky et al. (2017), which is equivalent to minimization of the earth mover's distance between the empirical (here, inferred from VB) and generated data distributions. That is, we minimized

$$\max_{w\in\mathcal{W}} \mathbb{E}_{x\sim p_{\mathcal{D}}}[D_w(x)] - \mathbb{E}_{z\sim p(z)}[D_w(G(z))] \tag{10}$$

subject to the restriction that $\mathcal{W}$, the space of GAN parameters, is compact. Both networks take s as input in addition to their normal inputs (noise for G and data for D). In addition, while most previous applications of GANs have focused on large, well-studied image data sets, in our case, the assumptions (6) and (8) are helpful, since even though the distribution $p(\mu, \lambda|s)$ is potentially quite complex, the choice to avoid conditioning on previous g and to model $V(g|s)$ as a smooth function of s allows us to reuse the same generator network not only on each bout of play, but at each time step within that bout. Thus, even for a modest number of trials and time points, we benefit from a multiplicative increase in training data to the GAN. Moreover, this is the first work of which we are aware that conditions the GAN on a *continuous* state variable, rather than a discrete class label, a novel application of the conditional GAN idea.

3. We have not yet described how we obtain these samples, since g and thus μ and λ are unobserved. We do so in Section 4.2.

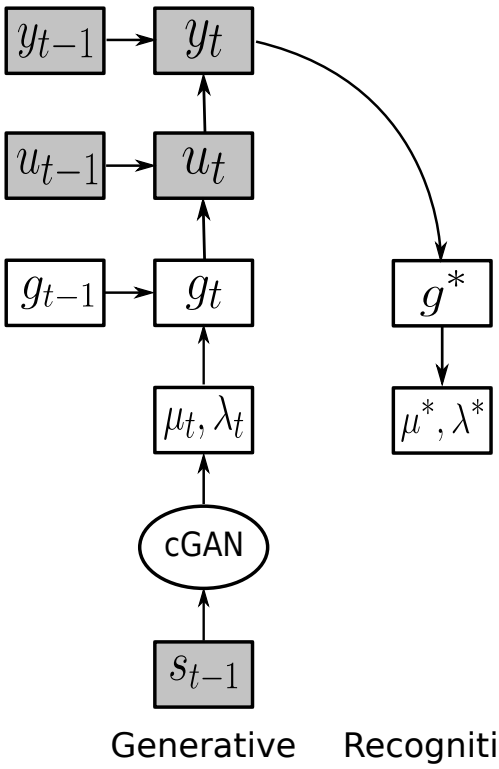

Figure 1: Model architecture. Observable trajectories y_t are generated sequentially from control signals u_t, which in turn derive from goals g_t. In the generative model, parameters of the equations governing goal evolution are drawn from the cGAN, which is conditioned on state variables s_t. The recognition model takes in y_t and returns posterior samples of the goals and cGAN parameters.

4. Full model and inference

4.1 Full generative model

Putting together the assumptions of the previous sections, we can now present a full generative model for our data (Figure 1). A given trial can be generated timepoint-by-timepoint as follows:

4.2 Approximate Bayesian inference

Given the observed system trajectory y_t (equivalently, the observed control signal u_t), we need to make inferences about the underlying goal trajectory g_t in order to obtain training data for the GAN. In general, full Bayesian inference is intractable, but we employ a variational Bayes (VB) approach Beal (2003); Wainwright et al. (2008) that approximates this procedure. In brief, VB attempts to minimize the Kullback-Leibler divergence between a known generative model for which inference is intractable, $p(\mathcal{D}, z)$, and an approximating family of posterior distributions, $q(z)$. This is equivalent to maximizing an evidence lower bound (ELBO) given by

$$\mathcal{L} = \mathbb{E}_{q(z)}[\log p(\mathcal{D}, z)] - \mathcal{H}[q(z)] \tag{11}$$

Algorithm 1 Generative model for trajectory data

Input: σ, ϵ, L, $v_{\max}$
  Initialize y_0, s_0, u_0
  $g_0 \sim G_{g_0}$ {Sample from GAN}
  for $t = 1 \dots T$ **do**
    $(\mu_t, \lambda_t) \sim G_{(\mu,\lambda)}(s_{t-1})$ {Sample from cGAN}
    $g_t \sim \mathcal{N}\left(\frac{g_{t-1}+\lambda_t \mu_t}{1+\lambda_t}, \frac{\sigma^2}{1+\lambda_t}\right)$ {Sample new goal}
    $e_t \leftarrow g_t - y_{t-1}$ {Control error}
    $u_t \sim \mathcal{N}(u_{t-1} + L * e_t, \epsilon^2)$ {PID control}
    $y_t \leftarrow y_{t-1} + v_{\max} \tanh(u_t)$ {Update position}
    $s_t \leftarrow (y_t, y_t - y_{t-1}, \dots)$ {Update state}
  end for

with $\mathcal{H}[q(z)]$ the entropy of the approximating posterior. That is, inference is transformed into an optimization problem in the parameters of the approximate posterior $q(z)$, amenable to solution by gradient ascent. In our model, we make use of recently developed "black box" methods Ranganath et al. (2014); Kucukelbir et al. (2015); Rezende et al. (2014); Kingma and Welling (2013) in which the gradients of the ELBO are replaced with stochastic approximations derived by sampling from $q(z)$, avoiding often difficult computations of the expectation in (11). Thus our only requirement for $q(z)$ is that it be straightforward to sample from.

In our case, we begin with the generative model specified by (4) and detailed in Algorithm 1. For the approximate posterior, we assume a factorized form $q = q(\mu, \lambda|g)q(g|y)$[4] modeled as follows: $q(g|y)$ is the variational latent dynamical system (VLDS) model of Archer et al. (2015); Gao et al. (2016); and $q(\mu_t, \lambda_t|g_t, g_{t-1})$ is a gaussian with mean and covariance parametrized by neural networks. As in Rezende et al. (2014); Kingma and Welling (2013), these samples are then used to update *both* the parameters of the generative model (L, ϵ, σ) and the parameters of the approximate posteriors (q_u, q_g, $q_{(\mu,\lambda)}$) via gradient ascent.

Algorithm 2 Variational inference on goal trajectories

Input: : data $\mathcal{D} = \{y, u\}$, tolerance ε
  while $\Delta\mathcal{L} > \varepsilon$ **do**
    Draw bout $(y^{(i)}, u^{(i)}) \sim \mathcal{D}$
    Draw $g^* \sim \text{VLDS}(g|y^{(i)})$
    Draw $(\mu^*, \lambda^*) \sim \text{NN}(\mu, \lambda|g^*)$
    `TrainVB`(g^*, μ^*, λ^*)
  end while
Output: : Generative examples $\widetilde{\mathcal{D}} = \{(\mu, \lambda, s)\}$

5. Experiments

5.1 Data

Our data consist of 10 sessions ($\approx$ 7000 bouts) of a simple two-player "penalty shot" video game played by pairs of rhesus macaques in a neuroscience experiment. Each player used a joystick to control either a "puck" or "ball" able to move in either the x or y direction or a "goalie" only able to move in the y direction (Figure 2). The objective of the player controlling the puck was to move it across a goal line at the right-hand side of the screen, while the objective of the goalie was to intercept the puck. Each bout of play ended when one of

4. We do not explicitly model $q(u_t|g, y)$, though $p(u_t|g, y)$ is normal at each time step. Since $\epsilon \ll 1$, this is effectively a delta function.

Algorithm 3 GAN training

Input: : data $\widetilde{\mathcal{D}} = \{(\mu, \lambda, s)\}$, tolerance ε
  while $\Delta\mathcal{L}_{\mathcal{D}} > \varepsilon$ **do**
    for $t = 1...K$ **do**
      Draw B "real" samples and conditions $\sim \widetilde{\mathcal{D}}$
      Generate B samples $(\mu^*, \lambda^*) \sim G(s_{t-1})$
      `TrainDiscriminator`$((\mu^*, \lambda^*, \mu, \lambda, s))$
    end for
    Generate B samples $(\mu^*, \lambda^*) \sim G(s_{t-1})$
    `TrainGenerator`$((\mu^*, \lambda^*, s))$
  end while

these two outcomes obtained. Data consisted of onscreen positions of both avatars at regularly sampled time points $y_t = (x_{puck}, y_{puck}, y_{goalie})$ for each bout of play (Figure 3). We normalize each coordinate such that $(y_t)_i \in [-1, 1]$, and we condition on state information $s_t \equiv (y_t, y_t - y_{t-1})$.

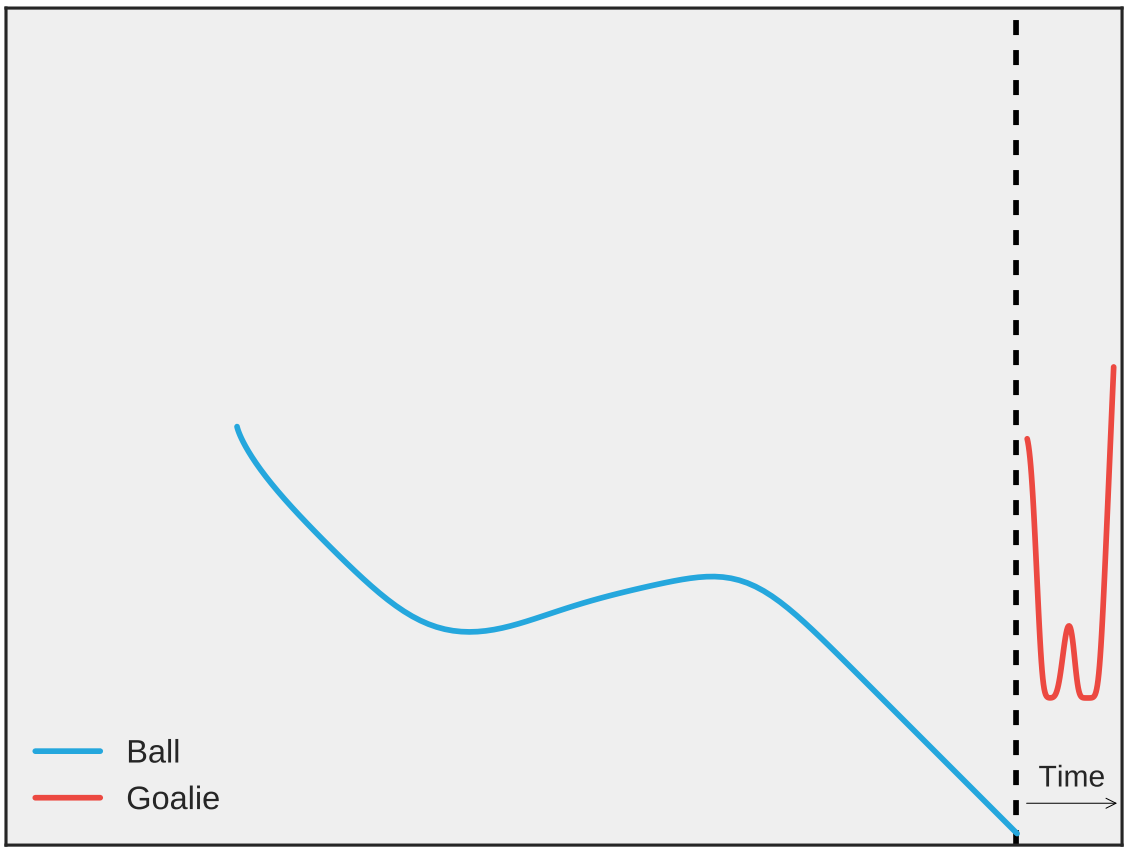

Figure 2: Game play. Illustration of onscreen play for the two-player game. The ball (blue) is free to move in two dimensions, while the goalie (red) moves only up or down. Axes are x and y screen directions. Solid lines indicate player trajectories for a single trial. The goalie's trajectory has been stretched in along the x axis to indicate its movement through time. For animations of game play, see Supplement.

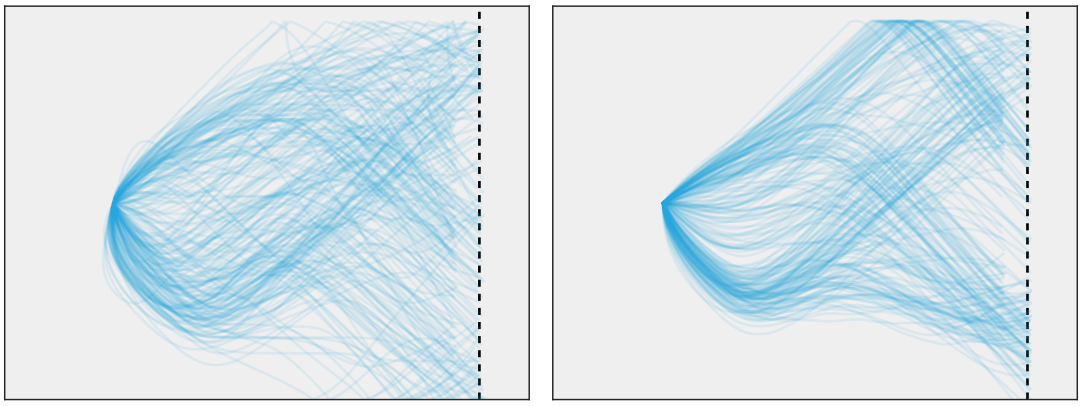

(a) Sample of 250 ball trajectories from training data. (b) Sample of 250 generated ball trajectories from Algorithm 1

Figure 3: Comparison of real and generated ball trajectories.

5.2 Training

Training consisted of two independent procedures: training of the VB posterior (Algorithm 2) and GAN training (Algorithm 3). For the posterior on g we used a three-layer neural network with 25 hidden units for each component of the mean and a network of the same structure for each nonzero entry of the Cholesky factor of the covariance as in Archer et al. (2015). The posterior $q(\mu, \lambda|g)$ was multivariate normal with each mean component parameterized by a neural network with three layers of 25 hidden units, and the Cholesky factor of a covariance matrix (with each diagonal entry constrained with a softplus) parametrized by a neural network with the same structure. We achieved best performance by setting $\sigma \sim 10^{-3}$ and including a regularization penalty for ϵ. Finally, to deal with nonidentifiability in the goal states arising from the sigmoid function applied to control, we incorporated a hinge loss on goals located outside the visible screen.

We trained using a single-trial black box approximation to the evidence lower bound Ranganath et al. (2014); Kucukelbir et al. (2015), utilizing ADAM Kingma and Ba (2014) with a learning rate of 10^{-3} to perform gradient ascent on both the hyperparameters of the generative model and parameters of the posterior approximation. The model was trained until a smoothed version of the evidence lower bound reached convergence. On the final iteration of VB training, we saved samples $\widetilde{\mathcal{D}} = \{(\mu, \lambda, s)\}$ to be used in training the GAN.

We perform GAN training using these samples collected from multiple trials and timepoints during the last phase of VB training. As in Arjovsky et al. (2017), we restricted the weights $w \in [-0.01, 0.01]$ but achieved increased performance by allowing the (unconstrained) batch normalization parameters for each layer to take values in $[-5, 5]$. For the discriminator (D), we used multilayer perceptrons (MLPs) with 4 layers, 35 hidden units, and leaky rectified linear units as our nonlinearity. We also used batch normalization before the nonlinearity in the inner layers (not input and output). For the generator (G), we sample 15 dimensions of noise from a uniform distribution as input to an MLP with the same structure as D. In addition to 3 training steps of D for each iteration of G, we also used a 25-epoch pre-training segment where we used 100 training iterations of D for each iteration of G. This allows the discriminator to approach an optimal level and provides a better objective for the beginning of training. We trained by using minibatches of 5000 samples, using RMSprop with learning rate $5e^{-5}$ to adjust parameters, until the WGAN objective no longer decreased.

Both networks take s as input in addition to their normal inputs (noise for G and data for D). Importantly, we found that adding additional gaussian noise to these states and annealing it towards 0 during training greatly improved our performance (Figure 4).

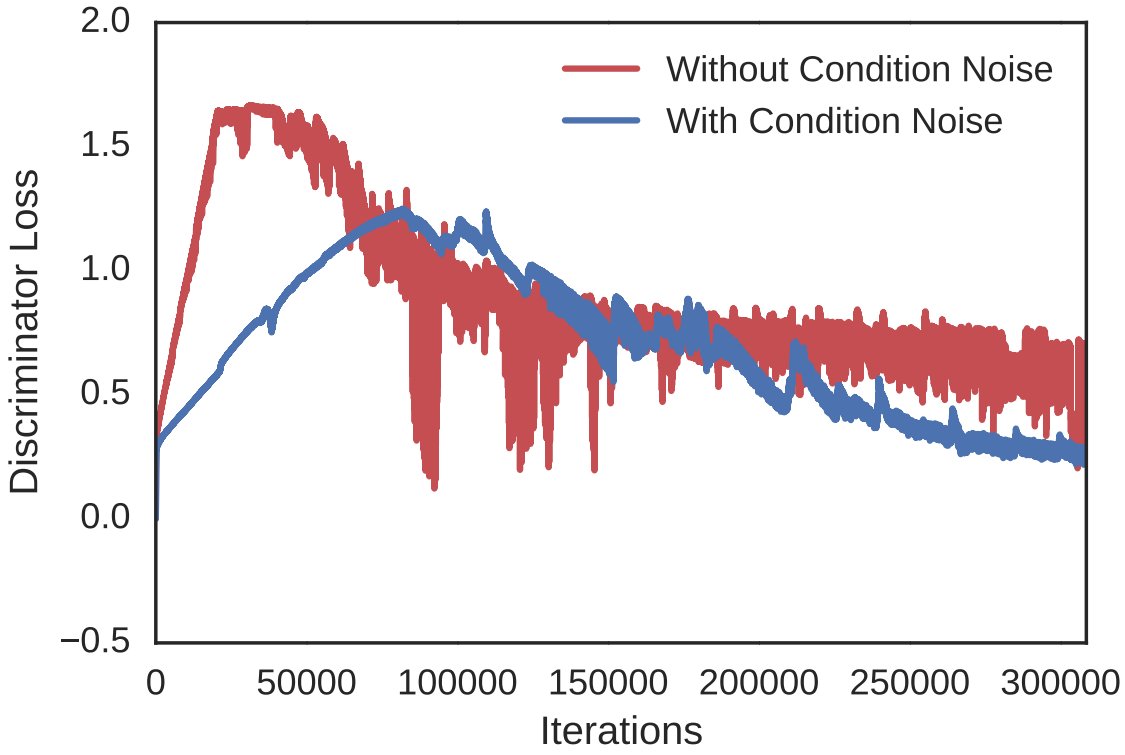

Figure 4: GAN training benefits from added condition noise. After several initial epochs of discriminator training, the WGAN objective decreases toward 0. Adding additional noise to the states s_t on which the GAN was conditioned resulted in a longer initialization time but a lower final optimization objective and qualitatively better samples.

Finally, for generating novel trials, we fit a separate WGAN to generate initial goals g_0, since these were highly multimodal (see Supplement). The D and G networks both had 4 layers with 20 hidden units and leaky rectifier nonlinearities. The generator used 5 dimensions of uniformly distributed noise as input. We found that adding instance noise to the input of the discriminator (annealed to 0 over epochs) vastly sped up our training. In this case we trained G and D equally after the pre-training stage where D is trained 100 iterations for each of 50 iterations of G. We used a batch size of 500 and learning rate of $5e^{-5}$.

5.3 Results

Our model is successfully able to capture the wide variability of player trajectories exhibited by player data (Figure 3). Moreover, the inferred value function $V(g|s)$ accurately captures the multimodal nature of player goal behavior (Figure 5): at a given moment in time, players are overwhelmingly likely to choose goals that remain along their current trajectory, but substantial probability mass is also placed on the possibility of an abrupt change in goals at points where this is strategically advantageous (see Supplement). Moreover, the GAN is able to generate these bifurcations when conditioned on a sparsely-sampled six-dimensional state space. As Figure 4 shows, this is significantly aided by the injection of additional noise to the s_t variables early in training. This may function similarly to instance noise in increasing the overlap between the data and generative distributions Sønderby et al. (2016), though WGANs should not benefit from instance noise

(and our G_{g_0} GAN does not) Arjovsky et al. (2017). More likely, this additional noise serves as a regularizer for the generative network, since even with several thousand trajectories of ≈ 100 time points each, the state space remains sparsely sampled.

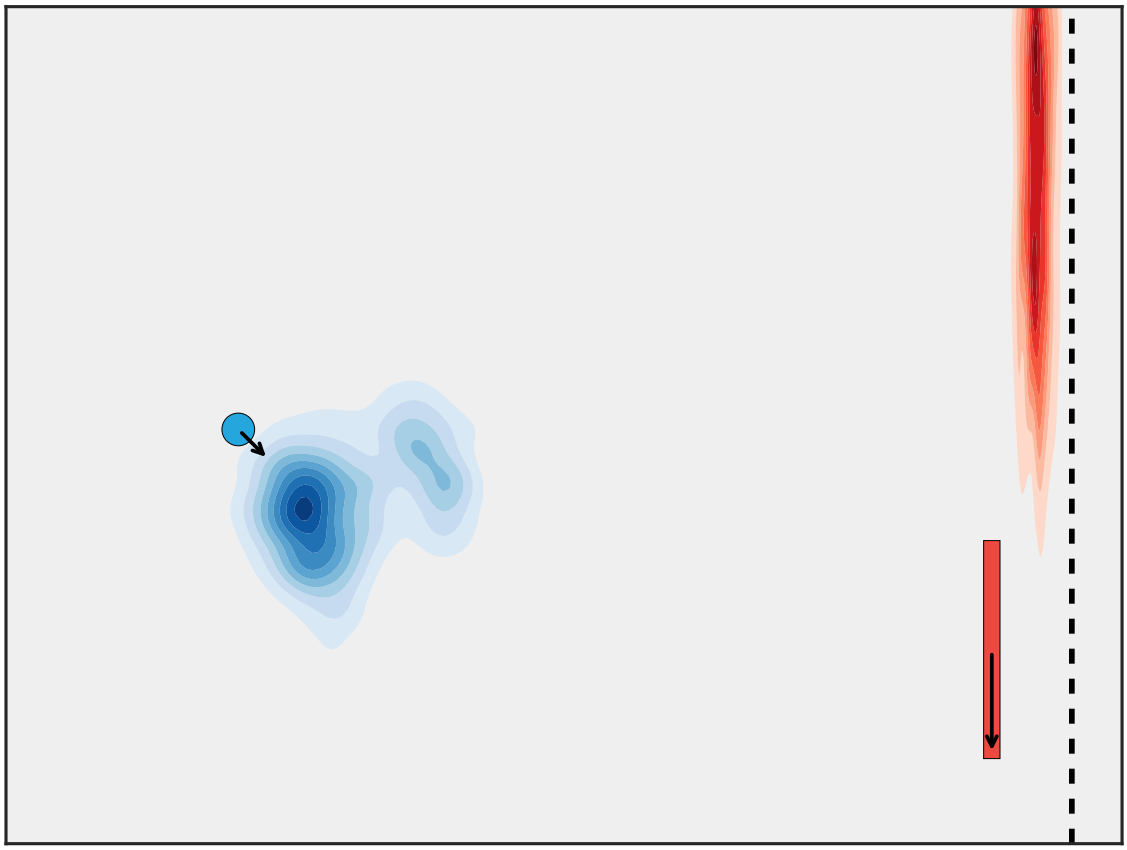

Figure 5: Value as a function of goal and state. The model correctly produces value functions for both players that exhibit the observed multimodality of real trajectories. Here, for given positions and velocities, the ball is likely to continue in its current direction, but retains some probability of shifting goals upward. Likewise, the goalie's value function exhibits suggests a bimodal distribution of goals at the next time point. Colors are as above. Heatmaps indicate kernel density estimates of e^V (see Supplement). For an animation of the evolution of the value function in time, see Supplement.

6. Related Work

Multiple related studies have made use of the idea of an energy or cost-to-go function in relation to GANs. The closest in spirit to this work is Finn et al. (2016), which uses a GAN to perform inverse reinforcement learning by learning a cost function over trajectories given optimal samples from a demonstrator. Our model differs in two key aspects: First, Finn et al. (2016) assume that the density for the generative distribution of the trajectory, $q(\tau)$, is calculable, which results in a closed-form expression for the optimal discriminator. We instead make the assumptions (6) and (7), which do not require the GAN itself to have a computationally tractable density. Second, Finn et al. (2016) approximate the cost function for entire trajectories, while we

focus on V, the cost of changing goals. By reusing the same GAN at every time step of every trajectory, we dramatically increase our training data for the same number of demonstrations, with the added cost of needing to condition on the current system state s.

Also related is the approach of Ho and Ermon (2016), which avoids value functions entirely, modeling the policy directly using a generative approach. Here, the difference is one of focus: we are interested explicitly in the function $V(g|s)$ as a means for understanding the rationale behind observed behavior, not only generating realistic actions to match a demonstrator.

Finally, several other papers, notably Kim and Bengio (2016); Zhao et al. (2016) use the idea of an an energy function to propose an energy-based GAN (EBGAN), which replaces the classification problem faced by the discriminator (is this sample real or synthetic) with a regression problem (minimize the energy of real samples while maximizing that of synthetic ones). Our model differs both in that our energy function is related to the *generator* and that we use a separate method Arjovsky et al. (2017) for GAN training.

7. Conclusion

We have proposed a model of inverse reinforcement learning based on inferring latent trajectories for goal states through time. These goal states give rise to observed states via a control model, and their evolution is governed by a Gaussian process determined by a dynamic value function. The model combines two leading generative approaches: a variational autoencoder (for inferring posteriors over goals given observations) and a generative adversarial network (for approximating V as a mixture of Gaussians). In addition, we make use of the recently developed WGAN model Arjovsky et al. (2017) to stabilize and assess training. To our knowledge, this is among the first applications of GANs to time series modeling and the only example we know of training GANs conditioned on continuous rather than discrete variables.

We have shown that such a model is able to reproduce the noisy and heterogenous behavior of real agents engaged in a competitive video game, and that by modeling the value function explicitly, we are able to easily visualize the multimodal distribution over future strategies for any instantaneous configuration of the system. This constitutes a significant improvement over conventional models in psychology, animal behavior, and related fields that seek to model the decisions of agents in dynamic environments, as it dramatically increases model flexibility while retaining the interpretability of a cost function over latent goals. For our example data, the inferred goal states, as well as their associated value functions, give important clues as to the underlying strategies employed by players, as well as offering potential neural correlates of the decision process. Moreover, this and similar approaches suggest new means of training more realistic artificial opponents for games by synthesizing strategies based on human demonstrators.

Acknowledgements

We would like to thank Michael Platt for sharing the data used in this work and Caroline Drucker for helpful discussions. This work was funded by NIH grants R01-MH109728 (PI: Platt; JP Co-Investigator) and a BD2K career development award (K01-ES025442) to JP.

References

Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1. ACM, 2004.

Evan Archer, Il Memming Park, Lars Buesing, John Cunningham, and Liam Paninski. Black box variational inference for state space models. 23 November 2015.

Martin Arjovsky, Soumith Chintala, and Lon Bottou. Wasserstein gan, 2017.

Matthew James Beal. *Variational algorithms for approximate Bayesian inference*. University of London United Kingdom, 2003.

Colin Camerer. *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press, 2003.

Chelsea Finn, Paul Christiano, Pieter Abbeel, and Sergey Levine. A connection between generative adversarial networks, inverse reinforcement learning, and Energy-Based models. 11 November 2016.

Yuanjun Gao, Evan Archer, Liam Paninski, and John P Cunningham. Linear dynamical neural population models through nonlinear embeddings. 26 May 2016.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*, pages 4565–4573, 2016.

Taesup Kim and Yoshua Bengio. Deep directed generative models with energy-based probability estimation. *arXiv preprint arXiv:1606.03439*, 2016.

Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. 22 December 2014.

Diederik P Kingma and Max Welling. Auto-Encoding variational bayes. 20 December 2013.

Alp Kucukelbir, Rajesh Ranganath, Andrew Gelman, and David Blei. Automatic variational inference in stan. In *Advances in neural information processing systems*, pages 568–576, 2015.

Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

Andrew Y Ng, Stuart J Russell, et al. Algorithms for inverse reinforcement learning. In *Icml*, pages 663–670, 2000.

Jooyoung Park and Irwin W Sandberg. Universal approximation using radial-basis-function networks. *Neural computation*, 3(2):246–257, 1991.

Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.

Rajesh Ranganath, Sean Gerrish, and David M Blei. Black box variational inference. In *AISTATS*, pages 814–822, 2014.

Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.

David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.

Casper Kaae Sønderby, Jose Caballero, Lucas Theis, Wenzhe Shi, and Ferenc Huszár. Amortised MAP inference for image super-resolution. *CoRR*, abs/1610.04490, 2016. URL `http://arxiv.org/abs/1610.04490`.

Martin J Wainwright, Michael I Jordan, et al. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, 2008.

Junbo Zhao, Michael Mathieu, and Yann LeCun. Energy-based generative adversarial network. *arXiv preprint arXiv:1609.03126*, 2016.

Supplement

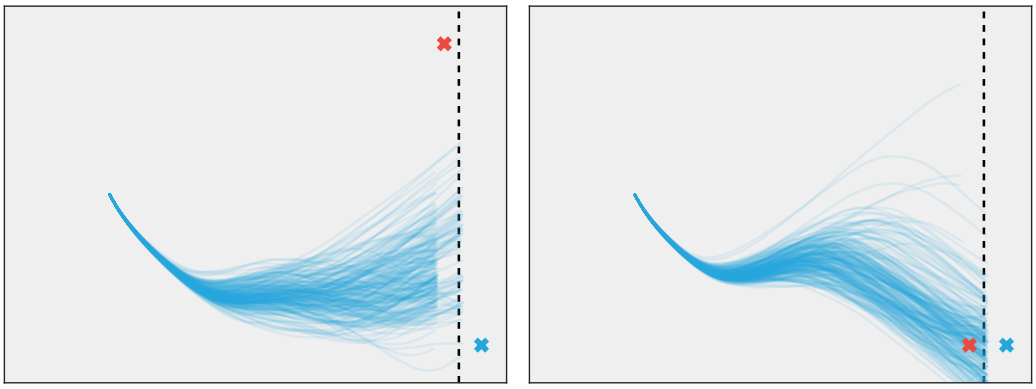

(a) Generated ball trajectories with initial goal states for ball in the bottom right corner and the goalie on the top of the screen. Note that the ball generally reaches a decision point late in trial in response to the goalie's downward motion in response to the ball's initial goal.

(b) Generated ball trajectories with initial goal states for ball in the bottom right corner and the goalie also on the bottom. Here, the ball generally shifts goal earlier because the goalie is initially headed downward, creating an earlier decision point.

Figure 6: Effects of initial goal state.

Approximating $V(g)$

Our model only encodes the value function $V(g|s)$ indirectly through a GAN that produces samples (μ, λ) conditional on s. However, for purposes of visualization and inference, it may useful to approximate V itself. From (7)

$$e^{V(g|s)} \propto \int p(\mu, \lambda|s) \sqrt{\frac{\beta\lambda}{2\pi}} e^{-\frac{\beta\lambda}{2}\cdot(g-\mu)^2} d\mu\, d\lambda \tag{12}$$

which we can approximate by assuming the existence of a radial basis function expansion for $p(\mu, \lambda|s)$:

$$p(\mu, \lambda|s) \approx \sum_{i=1}^{N} w_i K_\mu(\|\mu - \mu_i(s)\|/\kappa) K_\lambda(\|\lambda - \lambda_i(s)\|/\omega) \tag{13}$$

where we have assumed a normalized factorized basis function set in μ and λ.

In addition, in the limit of large β, we have

$$\sqrt{\frac{\beta\lambda}{2\pi}} e^{-\frac{\beta\lambda}{2}\cdot(g-\mu)^2} \to \delta(g-\mu) \tag{14}$$

so (12) becomes

$$e^{V(g|s)} \propto \sum_{i=1}^{N} w_i K_\mu(\|g - \mu_i(s)\|/\kappa) \tag{15}$$

It is known that in the limit of large N, the approximation (13) is always valid Park and Sandberg (1991). In practice, to avoid fitting the μ_i and w_i, we plot e^V via an importance sampling approach: draw N samples (μ_i, λ_i) from the GAN conditioned on some s and perform a kernel density estimate using a Gaussian basis function with bandwidth κ.

Initial Goal State, g_0

Based on the raw data (Figure 3a), it is clear that the players in our task had very distinct goals at the beginning of each trial. We found it was important to model the initial goal state with a GAN in order to capture the full variety of trajectories. Figure 7 demonstrates our success at replicating this distribution.

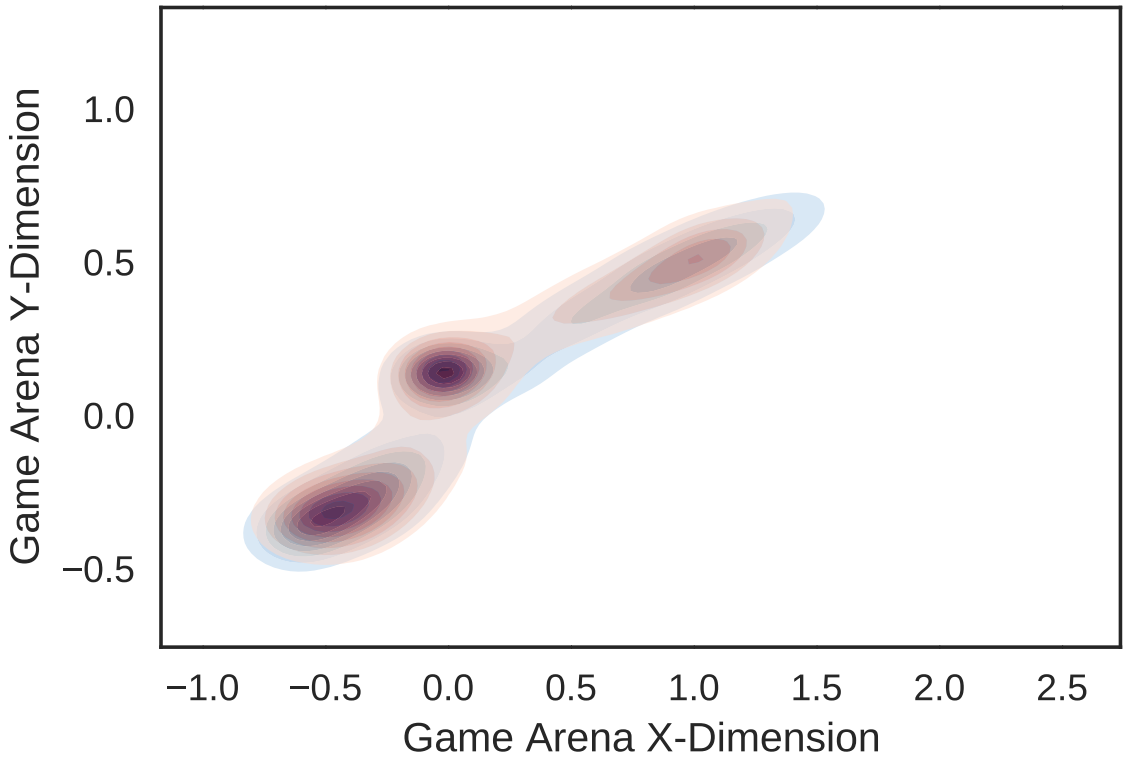

Figure 7: Density plot of inferred (blue) and generated (red) initial goal states, g_0, for the ball in our task. Our generative network is able to capture the multimodality, as well as the relative proportions between modes, in the inferred goal states.

In order to demonstrate the effect of initial goal state on trial outcome, as well as gain insight into the interaction between agents in our model, we generated trials with manually fixed initial goal states. These can be seen in Figure 6. Red Xs indicate the goalie's initial goal, and blue Xs indicate the ball's initial goal. Towards the beginning of the trial, the agents both strongly follow their initial goal, but as the trial progresses, they become more aware of their opponent. In the case of 6a, the goalie (not plotted for simplicity) guesses wrong initially and quickly turns around. At this point, the ball stops heading in that direction and decides to go up or down. In the case of 6b, however, the ball diverges from its initial goal more quickly because the goalie guesses correctly. At this point, the ball either tries to feint up and go down. These types of behaviors demonstrate the ability of our model to accurately capture the interaction inherent in the task, as opposed to a model that simply replicates trajectories.

Animations

In order to visualize the dynamics of our model in real time, we generated animations of individual trials (both real and generated from our model) with different visualizations.

- **Animation S1:** A real trial, played by monkeys. (link)
- **Animation S2:** The same real trial, with goal states inferred by our recognition model at each time point plotted with empty colored shapes corresponding to their agent. (link)
- **Animation S3:** The same real trial, with a value function estimate, obtained through samples from our cGAN, plotted at each time point. (link)
- **Animation S4:** A trial generated from our model. (link)
- **Animation S5:** The same generated trial, with corresponding goal states from our generative model. (link)
- **Animation S6:** The same generated trial, with a value function estimate, obtained through samples from our cGAN, plotted at each time point. (link)